	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments
1	BRS	L1	32	(register near3 file) SAME read adj port SAME write adj port SAME concurrent$5	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:15	
2	BRS	L2	20	1 and path	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:16	
3	BRS	L3	5	2 and ALU and multiplier	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:16	
4	BRS	L4	312	(register near3 file) SAME read adj port SAME write adj port and concurrent$5	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:15	
5	BRS	L5	82	4 and ALU and multiplier	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:16	
6	BRS	L6	71	5 and path	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:16	
7	BRS	L7	66	6 not 3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:16	

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments
8	BRS	L8	64	7 and first and second	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:17	
9	BRS	L9	45	7 and ((first or second) near6 port$3)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:18	
10	BRS	L10	38	9 and data adj path	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:18	
11	BRS	L11	38	10 and execution	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 12:18	
12	BRS	L12	38	11 and bus$3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 13:27	
13	BRS	L13	8	12 and cluster$5	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/09/17 13:27	

Terms used
VLIW and **register** and **read** and **write** and **port** and **data** and **path** and **concurrently** and **first** and **second** and **execution** an **file** and **read port** and **write port**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search
Try this search in The ACM Guid

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevan

1 Banked multiported register files for high-frequency superscalar microprocessors
Jessica H. Tseng, Krste Asanović
May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual internation symposium on Computer architecture**, Volume 31 Issue 2
Full text available: pdf(142.29 KB) Additional Information: full citation, abstract, references, citings

Multiported register files are a critical component of high-performance superscalar microprocessors. Conventio multiported structures can consume significant power and die area. We examine the designs of banked multip register files that employ multiple interleaved banks of fewer ported register cells to reduce power and area. B register files designs have been shown to provide sufficient bandwidth for a superscalar machine, but previou complex control structures that w ...

2 Superscalar microarchitecture: Register write specialization register read specialization: a path to comp effective wide-issue superscalar processors
André Seznec, Eric Toullec, Olivier Rochecouste
November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitectu**
Full text available: pdf(1.16 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

With the continuous shrinking of transistor size, processor designers are facing new difficulties to achieve high frequency. The register file read time, the wake up and selection logic traversal delay and the bypass network delay with also their respective power consumptions constitute major difficulties for the design of wide issue s processors.In this paper, we show that transgressing a rule, that has so far been applied in the design of all t superscalar processors, ...

3 The white dwarf: a high-performance application-specific processor
A. Wolfe, M. Breternitz, C. Stephens, A. L. Ting, D. B. Kirk, R. P. Bianchini, J. P. Shen
May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual Internation Symposium on Computer architecture**, Volume 16 Issue 2
Full text available: pdf(1.40 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents the design and implementation of a high-performance special-purpose processor, called T Dwarf, for accelerating finite element analysis algorithms. The White Dwarf CPU contains two Am29325 32-bi point processors and one Am29332 32-bit ALU, and employs a wide-instruction word architecture in which the algorithm is directly implemented in microcode. The entire system is VME-bus compatible and interfaces with host. The syste ...

4 Dynamically scheduled VLIW processors
B. Ramakrishna Rau
December 1993 **Proceedings of the 26th annual international symposium on Microarchitecture**

Full text available: pdf(1.64 MB) Additional Information: full citation, references, citings

Keywords: VLIW processors, dynamic scheduling, multiple operation issue, out-of-order execution, scoreboa

5 Data path issues in a highly concurrent machine
Augustus K. Uht, Darin B. Johnson
December 1992 **ACM SIGMICRO Newsletter , Proceedings of the 25th annual international symposium o Microarchitecture**, Volume 23 Issue 1-2
Full text available: pdf(492.18 KB) Additional Information: full citation, references, citings, index terms

6 Superscalar architectures: Reducing the complexity of the register file in dynamic superscalar processo
Rajeev Balasubramonian, Sandhya Dwarkadas, David H. Albonesi
December 2001 **Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitectu**
Full text available: pdf(1.34 MB) Publisher Site Additional Information: full citation, abstract, references, citings

Dynamic superscalar processors execute multiple instructions out-of-order by looking for independent operati large window. The number of physical registers within the processor has a direct impact on the size of this wi in-flight instructions require a new physical register at dispatch. A large multi-ported register file helps improv instruction-level parallelism (ILP), but may have a detrimental effect on clock speed, especially in future wire- technologies. ...

7 Dynamic dead-instruction detection and elimination
J. Adam Butts, Guri Sohi
October 2002 **Proceedings of the 10th international conference on Architectural support for programm languages and operating systems**, Volume 36 , 37 , 30 Issue 5 , 10 , 5
Full text available: pdf(1.50 MB) Additional Information: full citation, abstract, references, citings

We observe a non-negligible fraction--3 to 16% in our benchmarks--of *dynamically dead instructions,* dynami instances that generate unused results. The majority of these instructions arise from static instructions that a useful results. We find that compiler optimization (specifically instruction scheduling) creates a significant por *partially dead* static instructions. We show that most of the dynamically instructions arise from a small set of

8 MOVE: a framework for high-performance processor design
Henk Corporaal, Hans (J.M.) Mulder
August 1991 **Proceedings of the 1991 ACM/IEEE conference on Supercomputing**
Full text available: pdf(1.04 MB) Additional Information: full citation, references, citings, index terms

9 Two-level hierarchical register file organization for VLIW processors
Javier Zalamea, Josep Llosa, Eduard Ayguadé, Mateo Valero
December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitectu**
Full text available: pdf(154.90 KB) ps(843.85 KB) Publisher Site Additional Information: full citation, references, citings, index terms

10 Microprocessor architecture: A scalable wide-issue clustered VLIW with a reconfigurable interconnect
Osvaldo Colavin, Davide Rizzo
October 2003 **Proceedings of the international conference on Compilers, architectures and synthesis fo embedded systems**
Full text available: pdf(365.26 KB) Additional Information: full citation, abstract, references, index terms

Clustered VLIW architectures have been widely adopted in modern embedded multimedia applications for the exploit high degrees of ILP with reasonable trade-off in complexity and silicon costs. Studies have however sh performance scaling for wide-issue machines. In this paper we describe the architecture of a clustered VLIW w runtime reconfigurable inter-cluster bus suitable to address such scalability problem. The architecture is aime loops acceleration thr ...

Keywords: IDCT, clustered VLIW, modulo scheduling, reconfigurable co-processor (RCP)

11 Processor microarchitecture II: Reducing register ports using delayed write-back queues and operand
Nam Sung Kim, Trevor Mudge
June 2003 **Proceedings of the 17th annual international conference on Supercomputing**

Full text available: pdf(381.44 KB) Additional Information: full citation, abstract, references, citings, index terms

In high-performance wide-issue microprocessors the access time, energy and area of the register file are ofte overall performance. This is because these pararmeters grow superlinearly as read and write ports are added wide-issue. This paper presents techniques to reduce the number of ports of a register file intended for a wide microprocessor without noticeably impacting its IPC. Our results show that it is possible to replace the 16 read file of an eig ...

Keywords: instruction level parallelism, low power, out-of-order processor, register file, write queue

12 Register file and memory system design: Reducing register ports for higher speed and lower energy
Il Park, Michael D. Powell, T. N. Vijaykumar
November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitectu**

Full text available: pdf(1.28 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

The key issues for register file design in high-performance processors are access time and energy. While prev has focused on reducing the number of registers, we propose to reduce the number of register ports through proposals, one for reads and the other for writes. For reads, we propose bypass hint to reduce register port re by avoiding unnecessary register file reads for cases where values are bypassed. Current processors are unab these unnecessary reads due ...

13 A flexible datapath allocation method for architectural synthesis
Kyumyung Choi, Steven P. Levitan
October 1999 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 4 Issue 4

Full text available: pdf(195.48 KB) Additional Information: full citation, abstract, references, citings, index terms

We present a robust datapath allocation method that is flexible enough to handle constraints imposed by a va target architectures. Key features of this method are its ability to handle accurate modeling of datapath units simultaneous optimization of direct objective functions. The proposed method consists of a new binding mode construction scheme and an optimization technique based on simulated annealing. To illustrate the flexibility method, two datapath allocation ...

Keywords: allocation and binding, high-level synthesis

14 Register file port requirements of transport triggered architectures
Jan Hoogerbrugge, Henk Corporaal
November 1994 **Proceedings of the 27th annual international symposium on Microarchitecture**

Full text available: pdf(533.24 KB) Additional Information: full citation, abstract, references, citings, index terms

Exploitation of large amounts of instruction level parallelism requires a large amount of connectivity between register file and the function units; this connectivity is expensive and increases the cycle time.This paper show new class of transport triggered architectures requires fewer ports on the shared register file than traditional triggered architectures. This is achieved by programming data-transports instead of operations.Experimen ...

15 Register file and memory system design: Dynamic addressing memory arrays with physical locality
Steven Hsu, Shih-Lien Lu, Shih-Chang Lai, Ram Krishnamurthy, Konrad Lai
November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitectu**

Full text available: pdf(907.32 KB) Publisher Site Additional Information: full citation, abstract, references, index terms

As pipeline width and depth grow to improve performance, memory arrays in microprocessors are growing in ports. Arrays will increase in physical size, which prolongs the access time due to wiring delay. In order to boo frequency, these memory arrays must take multiple cycles to complete an access. This delays the scheduling instructions and affects overall performance. This paper proposes a different circuit organization to enable fas accesses solely de ...

16 Efficient checker processor design
Saugata Chatterjee, Chris Weaver, Todd Austin
December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitectu**

Full text available: pdf(153.00 KB) ps(1.26 MB) Publisher Site Additional Information: full citation, references, citings, index terms

17 Low-power: Low-complexity reorder buffer architecture
Gurhan Kucuk, Dmitry Ponomarev, Kanad Ghose
June 2002 **Proceedings of the 16th international conference on Supercomputing**

Full text available: pdf(120.97 KB) Additional Information: full citation, abstract, references, citings, index terms

In some of today's superscalar processors (e.g.the Pentium III), the result repositories are implemented as th Buffer (ROB) slots. In such designs, the ROB is a complex multi-ported structure that occupies a significant po die area and dissipates a non-trivial fraction of the total chip power, as much as 27% according to some estim addition, an access to such ROB typically takes more than one cycle, impacting the IPC adversely.We propose complexity and low-powe ...

Keywords: low-complexity datapath, low-power design, reorder buffer

18 Performance comparison of ILP machines with cycle time evaluation
Tetsuya Hara, Hideki Ando, Chikako Nakanishi, Masao Nakaya
May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual internation symposium on Computer architecture**, Volume 24 Issue 2

Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, citings, index terms

Many studies have investigated performance improvement through exploiting instruction-level parallelism (ILP particular architecture. Unfortunately, these studies indicate performance improvement using the number of c are required to execute a program, but do not quantitatively estimate the penalty imposed on the cycle time architecture. Since the performance of a microprocessor must be measured by its execution time, a cycle tim is required as well as a cy ...

19 Processor coupling: integrating compile time and runtime scheduling for parallelism
Stephem W. Keckler, William J. Dally
April 1992 **ACM SIGARCH Computer Architecture News , Proceedings of the 19th annual internation symposium on Computer architecture**, Volume 20 Issue 2

Full text available: pdf(1.32 MB) Additional Information: full citation, abstract, references, citings, index terms

The technology to implement a single-chip node composed of 4 high-performance floating-point ALUs will be a 1995. This paper presents processor coupling, a mechanism for controlling multiple ALUs to exploit both instr and inter-thread parallelism, by using compile time and runtime scheduling. The compiler statically schedules threads to discover available intra-thread instruction-level parallelism. The runtime scheduling mechanism int threads, explo ...

20 Effective cluster assignment for modulo scheduling

h c g e cf c

Erik Nystrom, Alexandre E. Eichenberger
November 1998 **Proceedings of the 31st annual ACM/IEEE international symposium on Microarchitectu**

Full text available: pdf(1.72 MB) Additional Information: full citation, references, citings, index terms

Keywords: ILP, cluster architecture, cluster assignment, modulo scheduling

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next